

Beamery AI Explainability Statement

September 2024_v3

beamery.com

Private & Confidential - Do Not Share © Beamery Inc. All rights reserved



TRANSPARENCY IN ARTIFICIAL INTELLIGENCE (AI) IS ONE OF OUR CORE VALUES AT BEAMERY.

In this regard, we share the perspective of the Alan Turing Institute - that being transparent about the AI system itself and how it was developed, with respect to decisions made towards fairness, reliability, and accountability, is the best way to provide an understanding of the AI system, and ultimately instill trust¹.

This Explainability Statement outlines our AI processes in order to instill trust through understanding, with the additional aim of assisting our customers in fulfilling their obligations under applicable laws and/or due diligence requirements around transparency in AI in the procurement and deployment process.

We believe that HR should play a crucial role in the use and management of AI, as well as in ensuring that pro-innovation attitudes towards AI are matched by its safe and effective use. This guide is intended for those who may not have a background in coding AI systems, yet are accountable for making decisions about utilizing responsible AI solutions. To this end, this guide has been written within the context of not only the evolving technology landscape but the regulatory one.

RECENT GLOBAL REGULATORY CONTEXT FOR THE USE OF AI


In March 2024, the UK Department for Science, Innovation, and Technology published guidance on **Responsible AI in Recruitment**, available [here](#). This government guidance sets out processes and assurance mechanisms for employers to consider and implement before, during, and after AI procurement and deployment in the workplace.

The EU AI Act (the "Act") entered into force on August 2, 2024, and, looking ahead, most of its obligations will go into effect 24 months later, on August 2, 2026². The Act sets out a comprehensive set of transparency requirements that will apply to AI technology Providers and Deployers (as defined in the Act) and will co-exist with other transparency requirements already in place, such as the GDPR. As a company in the highly regulated European market, Beamery has always held regulatory compliance and human rights at the core of our business ethos, keeping regulations such as the GDPR and EU Employment Equality Directives front of mind when developing all of our services. We are also closely monitoring expected global AI legislation, such as the Colorado and California AI acts.

¹Information Commissioners Office (ICO) & The Alan Turing Institute. (2020). Explaining decisions made with AI. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/explaining-decisions-made-with-artificial-intelligence/>

²Note: the EU AI Office anticipates to provide guidance on High-Risk System no later than 18 months from the enactment date, so this information is from the prima facie reading of the Act. <https://artificialintelligenceact.eu/ai-act-implementation-next-steps/>

This Explainability Statement Demonstrates Beamery's AI in three parts:



PART A

How the AI in the Beamery suite of products works and how utilizing this technology helps our customers maximize human potential, aka its purpose;



PART B

Beamery's commitment (as the Provider and Distributor of AI) to the Act's principles, and the UK Responsible AI in Recruitment guidance, as well as our commitment to helping our customers (the Deployer of AI) prepare for the same, and our commitment to our customers' justified trust in Beamery as a supplier; and



PART C

How Beamery AI meets the pillars of the responsible AI procurement framework, so customers trust Beamery as an AI vendor.

Part A

Our Suite of Products and their Use of AI

Prior to procuring an AI system, an organization should develop a clear vision that outlines its desired **purpose** for using it. Understanding this purpose will also help ensure that the system is used to perform the task for which it was designed, and avoid any potential misuse.

We believe that data-driven decision making enabled by technology is the future of HR. If embedding fairness in the selection process is your organization's focus, then technology is the most reliable and scalable way to achieve that. This is where Beamery steps in. Beamery's AI-powered talent platform helps businesses identify, engage, recruit, develop, retain, and redeploy talent more effectively, by unlocking the skills and capabilities of their global workforce.

As we stand on the cusp of an AI Revolution, ensuring prosperity for all becomes increasingly crucial. Beamery contributes to this purpose by enabling employers to harness the vast potential of AI to utilize a Skills-Based Hiring approach for the benefit of both employers and candidates.

HOW AI IS DEPLOYED BY BEAMERY IN ITS OFFERING CAN BE SEPARATED INTO TWO ELEMENTS:



1. Beamery's self-developed AI models and Job Architecture³ that power certain features in the platform, such as matching applicants to vacancies and suggesting contacts for a job requisition; and



2. Generative AI features and functionality that is powered by Microsoft Azure's OpenAI LLM⁴, which has been fine-tuned using Beamery's robust HR taxonomy (capabilities, skills, roles, organizations, degrees, job families, etc) and how Beamery models them ("TalentGPT"). TalentGPT helps users get all the advantages of natural language interaction, with the enterprise compliance and security controls that are critically important in handling Personal Data and ensuring organizational compliance. TalentGPT is more efficient and accurate than ChatGPT alone, as it tailors the customer experience to align with the customer's unique job architecture.

³ A job architecture is a structured framework that defines and organizes roles within an organization, including their hierarchy and relationships to other positions.

⁴ Beamery partnered up with Microsoft to deliver enterprise-level LLM capabilities and as such, Beamery has its own private instance of Azure OpenAI, which does not feed any data or information back to Microsoft. More information on Azure OpenAI data privacy practices can be found here. <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy>

Beamery enables large enterprises to become skills-based organizations, by bringing together for the first time a connected skills-based AI ecosystem - which includes connecting customers to Beamery's "universal skills platform" as well as leveraging Beamery's talent management suite that brings together Job Architecture, CRM, and Talent Marketplace⁵.

Through these connected solutions, Beamery provides a unique, unified experience for internal and external resourcing, and helps organizations close key skill gaps, retain top performers, and improve DE&I efforts/outcomes. Beamery AI supports the user experience across the platform by enhancing recruiter and talent manager productivity, and their understanding of individuals within the system. It also unlocks new and exciting ways for individuals to think about their careers. Specific projects under these themes include smart data reconciliation within the CRM, enhanced skills inference, career pathing recommendation (given a specific career goal), learning and development optimization, fuzzy entity matching to account for various ways of stating the same role/skill, and ways to improve diversity and inclusion, via directed engagement such as event marketing towards specific under-represented communities.

Technologies such as AI and machine learning are powerful tools to help employers sift through extremely large amounts of data; for example, through functionality that matches individuals to vacancies that fit their skills and explains why they are a match in a matter of seconds. This enables TA Teams to interpret resumes better and suggest more suitable matches, based on the most relevant data, such as skills, job title, and location. In both cases, Beamery's AI helps to do this far more efficiently than traditional methods, allowing TA Teams to handle increased workloads, meet growing business demands, and enhance productivity without sacrificing quality. This, in turn, enables employers to reduce their time and costs to fill a vacancy.

Beamery's AI is beneficial not only for employers but can also be valuable for prospective candidates, applicants, and employees. For individuals, AI can lead to greater personalization and more curated recruitment experiences. For example, Beamery's AI (through Talent Portals) can surface higher-quality matches for open job requisitions, streamlining the job-seeking experience and highlighting internal and external opportunities with greater relevance. This helps to improve the overall user experience and enables more efficient job-hunting experiences for individuals.

Many studies⁶ have confirmed that when it comes to hiring decisions, unconscious bias is a significant factor in causing the unfair distribution of opportunities and decreasing diversity in the labor market. Overcoming unconscious bias by utilizing AI in your hiring can have a ripple effect of building an exceptional team that attracts exceptional candidates. This academic study⁷ states that there are "substantial gains in diversity from the use of AI in recruitment, both when isolating supply and demand effects and when integrating these effects together." A further encouraging statistic from this study⁷ says, "we find that adding AI to recruitment increases the representation of women at the 50th percentile of evaluated applicants by 30% and the 90th percentile of evaluated applicants by 160%."

By utilizing Beamery's transparent and explainable talent matching features, which match individuals to jobs based on objective and relevant criteria that focus on potential (such as relevant skills and job titles), rather than credentials/pedigree (such as school, degree, or company name), employers can mitigate many areas of unconscious bias that their recruiters or those involved in hiring may have⁸. Skills based hiring focuses on what a candidate "can do" vs "who they are". This objective approach both should and has been demonstrated to increase the talent pool from which to select candidates/applicants, as such, providing greater opportunities to undervalued and underrepresented talent. Current research suggests that if used carefully, AI can help avoid discrimination, and even raise the bar for human decision-making. When searching with Beamery Talent Match, the user is unable to search on any protected attributes, like gender or religion, to help prevent any direct discrimination.

⁵ a talent marketplace is a collection (could be a platform) of job descriptions/calibrations across geographies (demand side) and the collection of public profiles for people (supply). The raw data provide the base for workforce insights and analytics.

⁶ Giusta, M.D., & Bosworth, S. Bias and discrimination: what do we know?. Oxford Review of Economic Policy 36.4: 925-943. (2020).

⁷ Avery, Mallory and Leibbrandt, Andreas and Vecchi, Joseph, Does Artificial Intelligence Help or Hurt Gender Diversity? Evidence from Two Field Experiments on Recruitment in Tech (2024). CESifo Working Paper No. 10996, Available at SSRN: <https://ssrn.com/abstract=4764343> or <http://dx.doi.org/10.2139/ssrn.4764343>

⁸ <https://www.bcg.com/publications/2023/rise-of-skills-based-hiring>

Upskilling your workforce

In today's talent market – where critical skills are becoming increasingly scarce – companies recognize the importance of supporting employee skill development through structured programming. With the rapid advancement of AI (including GenAI), this trend has accelerated.

Research suggests that technological developments will require over 100 million workers to find new professions by 2030⁹. This scale of change will transform how people work and how businesses remain competitive to survive over the long term, and success hinges upon having a workforce with the right skills. Fortunately, the same advancements that contribute to the need for new skills also support the solution: upskilling. There has never been a stronger understanding of the connection between individuals and skills than we have today – and for much of that, we have AI to thank. Using AI organizations are able to draw correlations between knowledge and experiences to better understand what skills an employee likely has, and where there are gaps relative to expectations for a given role. This represents an evolution of thought and a step change for managers and employees alike.

Consider a manager who is new to a team; this sort of insight into teammates' skills can help inform decision-making around who to deploy and how, isolate specific areas for growth and contraction, and provide tailored recommendations for professional development. For employees, it means understanding developmental opportunities – including what is required to get to the next level – and seamlessly connecting employees to learning programs, in order to close any gaps.

By enabling your employees to harness the value of AI in their everyday tasks – by utilizing Beamery's AI technology (including TalentGPT) – you are enabling your employees to acquire the AI skills needed for this role and the next, and unlocking a world of opportunity. Further, this demonstration of how AI complements – not replaces – employees' work helps establish a level of comfort and confidence in the utility of AI, helping quell apprehension around its use in the workplace.

A comprehensive overview of Beamery's AI-powered capabilities is available [here](#). For more detailed information about our offerings, including an in-depth look at our AI, visit the Beamery AI Customer Hub or contact us directly to learn more.

⁹ <https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages>

Part B

Our Commitment to Responsible AI

OUR AI PRINCIPLES:

The EU AI Act has six principles that providers of AI systems should follow, and Beamery has adopted these principles:



1. Human Agency and Oversight



2. Technical Robustness and Safety



3. Privacy and Data Governance



4. Transparency



5. Diversity, Non-discrimination, and Fairness; and



6. Social and Environmental Well-Being.

Numerous organizations worldwide, including the Equal Employment Opportunity Commission (EEOC)¹⁰ and the United Nations System Chief Executives Board for Coordination¹¹, reflect similar principles regarding AI governance and utilization.

¹⁰ <https://www.eeoc.gov/laws/guidance/select-issues-assessing-adverse-impact-software-algorithms-and-artificial>

¹¹ <https://unsceb.org/sites/default/files/2024-05/United%20Nations%20System%20White%20Paper%20on%20AI%20Governance.pdf>

1. Human Agency and Oversight:

AI systems should assist humans in decision-making, and humans should be able to override decisions made by the system.

We empower our human users! Our AI is supercharging everyday tasks in the talent lifecycle, not replacing them. Unlike your typical ATS and some other CRMs, Beamery's AI does not make employment or workplace decisions or screen out individuals; a human (from your company) is always the decision maker in every step of the recruitment process within Beamery.

Human Governance: Beamery's AI models and features are not meant to replace humans; instead, they give relevant information to human decision makers to make more informed decisions. Utilizing an AI-powered, skills-based hiring process creates transparent and explainable outcomes that can demonstrate a candidate's suitability for the job based on objective criteria. This approach eliminates the unexplainable "black box" of the human mind, which makes it difficult, if not impossible to be able to explain why an individual was interviewed, hired, or did not progress in the process. Our system avoids the peril of a particular human reviewing a candidate strongly or poorly based on personal preferences that have nothing to do with job competency.

Human in the Loop: Every Beamery feature is designed to help inform our customers' talent-related decisions and requires a "human in the loop" to act on any output (i.e. decision to interview, hire, or not). To mitigate the risks of using AI, customers should ensure that qualified professionals review the AI output prior to making any next steps in the talent acquisition process. HR professionals can oversee AI outputs and decisions, helping to uphold ethical standards, maintain employee trust, and comply with legal and ethical HR standards. It is important to note that Beamery AI does not make any automated decisions; it does not screen/promote/reject candidates, nor does any step in the hiring process happen without a human making the decision.

Human Agency over the Candidate Experience: On the other side of the hiring process are the individuals who are searching for jobs. Studies¹² have shown that women are statistically less likely to apply to jobs they are qualified for than men. These differences can have a significant impact on the underrepresentation of protected groups in certain industries. To combat such disparities, Beamery provides, via its Candidate Portals, potential applicants with an objective analysis of their fit for a job, in order to encourage qualified applicants to apply. Additionally, by providing a personalized list of open jobs, ranked by compatibility, self-selection biases among potential applicants are further mitigated. This allows each job opportunity to reach the strongest potential applicants and increases the likelihood of a diverse applicant pool for recruitment.

AI can do better than humans: Unlike human bias, AI bias is testable and hence fixable. Beamery training data and AI models are both tested for bias; as such, the appearance of any bias can lead to corrective actions. Unlike the human unconscious, AI can be (and has been) trained to ignore demographic information such as gender, race, and age, and other traditional proxies for success. Our continuous monitoring of job architecture, candidate data, and predictions, will lead to an improved understanding of bias and further enhanced predictive models.

Further, utilizing objective criteria and not traditional proxies for "quality" (for ex: recommendations, prestigious education, well-known prior employers, years of experience, lack of gaps in employment, appearance (including digital presence), shared backgrounds), Beamery's AI should expand the suitable talent pool presented to a recruiter, and helps employers focus on evaluating the skills only. For example: when searching for a software engineer, many recruiters may choose to recruit only from MIT or Harvard or for employees who used to for Google or Tesla; the Beamery AI will instead suggest all candidates with relevant skills - derived from their work or educational experiences - even if they did not go to university or work in a top "credentialed" firm.

¹² <https://hbr.org/2014/08/why-women-dont-apply-for-jobs-unless-theyre-100-qualified>

2. Technical Robustness and Safety:

Providers and developers of AI systems should be designing AI systems to work well, be predictable, and be safe to use. AI providers should ensure that their systems comply with quality management systems.

To ensure the trustworthiness of our AI, we utilize robust testing and validation methods, and a proactive approach to identifying and mitigating risks. When an AI system makes decisions or predictions, “Model Monitoring” involves keeping an eye on how well it’s doing, ensuring it stays accurate, fair, and reliable over time. AI model monitoring ensures the AI isn’t making mistakes, becoming biased, or starting to perform poorly as time goes on. If something seems off, adjustments are made to keep it on track. Beamery employs model monitoring capabilities to assess the model’s prediction input data for feature “skew” and “drift” as explained below:



Training-serving “skew” occurs when the feature data distribution in production deviates from the feature data distribution used to train the model. If the original training data is available, you can enable skew detection to monitor your models for training-serving skew and Beamery fixes any skew if detected.



Prediction “drift” occurs when feature data distribution in production changes significantly over time. If the original training data isn’t available, you can enable drift detection to monitor the input data for changes over time and Beamery fixes any drift if detected.

We have robust security measures in place to protect data from attacks throughout all phases of the AI development and deployment lifecycle. We maintain industry-leading information security protections for all data stored on our systems (see further information on our information security practices at the Beamery Trust Center [here](#)). We adhere to the highest security and privacy practices to help anticipate and mitigate unintended harm and keep our products safe; this includes ISO 27001 and SOC 2 Type 2 certifications.

3. Privacy and Data Governance & Compliance:

AI system providers and developers should be designing AI systems with data privacy and data protection in mind. The datasets used to train AI systems should be properly governed. As with every element on the core Beamery platform, our AI has been developed with privacy by design and by default.

Personal Data Processing Compliance: Our platform enables our customers (the employers) to lawfully collect Personal Data during the hiring process and during employment in their role as a Controller (or the local law equivalent). Personal Data added to the Beamery systems is processed by Beamery as a Processor (or the local law equivalent) on behalf of our customers.

Further information about the data processed by Beamery can be found in our Data Protection Impact Assessment (DPIA) [here](#). This document can enable our customers to fulfill their obligations as “employers” or “data controllers” with regard to transparency and notice under certain laws.

Data Sets for Training: The data set used for training Beamery's AI was legally purchased from a third-party company, with warranties on the lawful collection of data, and consisted of millions of profiles (over one billion data points). After performing a Data Protection Impact Assessment (DPA), we stripped the data set of all data points that were not relevant or necessary for training (for example, all Personal Data used for identification: names, emails, addresses, etc). When training our models, we used these de-identified third-party datasets, which means our models are not explicitly aware of race, ethnicity, gender, nationality, sexual orientation, and political or religious beliefs.

Additionally, when we are validating our models against test data sets, non-Personal Data is used. When we match contacts in the CRM, our models never see the Personal Data of any contact, as the only data points the models need are job title, skills, organization size, industry, and seniority.

Employment Law Compliance: Recently, a former EEOC Commissioner stated that utilizing skills-based AI may be the most compliant form of hiring¹³, because an employer can demonstrate the use of objective criteria to an applicant or a regulator.

“If you're using AI to make employment decisions or to assist you in making employment decisions, there's a lot of clicks in there. And there's a lot of ways that you can use the AI to show, “Here are the skills we were looking for. Here's the job, here's the job description, here's the applicant pool, and here's what we use the machine learning on neutral characteristics that are related to the job, and here's the pool that came out of it. Okay? And here's why it was a business necessity of why we put these factors in to take that skills-based approach.”

And you have the record of that... I really think that this can really help not only with the EEOC's investigations, but help corporations have a much more transparent, documented way of how employment decisions are made, which has been lacking for a very long time. And also from an employee fairness perspective, if they know those qualifications that are going into the job, as opposed to, “Well, I just didn't get it because they didn't like me,” or, “I didn't get it because I'm this color,” versus, “No, this is the very transparent way we made the employment decisions.”

- EEOC's **Keith Sonderling** on job fairness in the age of AI." Managing the Future of Work, Harvard Business School, 20 Sep 2023

4. Transparency:

AI systems should be transparent. AI providers should provide clear information about the system's capabilities and limitations, as well as the data sources used to train it.

AI systems need to be transparent, with providers giving clear insight into:

- **The purpose of an AI system:** See Part A above (Our Suite of Products and their Use of AI).
- **The data processed by the system:** Beamery processes the data that is in your Beamery system; as further stated in Annex 1 of the DPA executed between the parties.
- **How the AI System works:** Beamery's AI features and functionalities utilize multiple models working together to achieve the desired output/results. See the diagrams in Figure 1 and Figure 2 below. It is important to note that the system has limitations as to "accuracy" (see more below in Section C. 2. a) as accuracy is difficult to measure in AI systems.
- **Decisions made by the system:** Beamery achieves **transparency with explainability** by helping users understand why we predict that someone will be a good fit for a role, using a few key explanation layers within the product. This can, in turn, allow customers to provide contacts with more objective explanations as to why they are a match for a job. Users can see each component's weight (or influence) in the star rating: our AI recommendations articulate the mix and weight of factors such as job title, skills, organization size, industry, and seniority. The Beamery Knowledge Graph maps an Organization to their respective size and industry to enable the matching functionality. This enables users to understand the reasons our AI considers a potential match between talent and vacancies. Users can understand what skills impacted the recommendations the most, also helping them to understand what skills did or did not have an impact on the recommendation. See 'Transparency with Explainability' below.
- **Where AI is in use:** We clearly indicate throughout our product where AI is used, so that users understand when AI is interacting with them. We employ a system of checks and balances to review user interactions and feedback to the platform, ensuring that the systems operate as they should.
- **Capabilities:** See Section 1 above.
- **Data Sources:** See Section 3 above (Privacy and Data Governance).

¹³ "The EEOC's Keith Sonderling on job fairness in the age of AI." Managing the Future of Work, Harvard Business School, 20 Sep 2023, <https://www.hbs.edu/managing-the-future-of-work/podcast/Pages/podcast-details.aspx?episode=1110268856>




Figure 1: Diagram depicting how Beamery's AI models work together.

5. Diversity, Non-discrimination, and Fairness:

Developers and providers of AI systems should be creating AI systems to avoid discrimination and bias and promote diversity. Providers should carefully examine data sources for bias and use proper measures to mitigate any potential biases.

AI has the potential to transform employment processes. So, fairness in AI is essential to address the risk of perpetuating potential social and unconscious biases and discrimination. This is why Beamery has prioritized fairness at every stage of development and testing our AI.

Data sources: We ensure our AI models are trained on a diverse range of data. The data set used by Beamery consisted of millions of profiles and over one billion data points, including profiles from different demographics, backgrounds, and scenarios to expose the models to a broader set of circumstances and reduce the risk of any bias. Further, to prevent amplifying any implicit biases, we have designed our products and services to avoid utilizing protected characteristics (such as race, ethnicity, gender, nationality, sexual orientation, and political or religious beliefs) by ensuring that these data points (as well as known proxies for such) are never included in the dataset for our training models. This means that the Beamery system can disregard any data points that are irrelevant for the successful execution of the job (e.g. gender, ethnicity, and age) and focus on those objective elements (such as skills), thus promoting diversity and inclusion within your organization. Utilizing a system of matching based on this objective, measurable criteria can reduce bias in your recruitment process. Our model outputs have been audited by a third party to confirm we are meeting the industry standards minimizing the risk of unintended bias.

Model Selection: From the very outset, Beamery was thinking about implementing fairness in our AI products. Our Entity Reconciliation Model considers different english gendered role titles, for example, pulling Barman and Barmaid into one entity: Bartender. Similarly, our Candidate Match Model doesn't consider characteristics such as race, ethnicity, gender, etc., instead focusing on elements that are relevant for the successful execution of the job: job title, skills, organization size, industry, and seniority.

Diverse Teams: Beamery has always emphasized the importance of interdisciplinary teams and diverse perspectives when creating its models. Diversity of thought is crucial to ensure that models are robust, fair, and understand a wide range of perspectives. The team working on our AI models consists of experts from various fields and diverse backgrounds, including machine learning, natural language processing, computer science, ethics, security, social sciences and legal experts.

This multidisciplinary approach is essential for tackling the multifaceted challenges associated with creating and refining the models at Beamery. Further, the Diversity in Beamery's AI development team brings a range of perspectives, experiences, and backgrounds to foster ideas and approaches that may have otherwise gone unnoticed. Beamery also actively collaborates with external researchers, solicits feedback from the wider community and customers, and is involved in various partnerships and initiatives that bring in different viewpoints.



Beamery's Bias Auditing Process:

- **Internal Testing:** As part of the model development, Beamery has a dedicated and rigorous testing phase where we internally audit machine learning predictions for equal parity, proportional parity, false positive parity, and false negative parity.
- **Third-party Auditing and Assurance (Warden AI):** Beamery is at the forefront of AI trust and safety. To provide external assurance that our AI is fair and trustworthy, Beamery engages an independent third-party AI auditor, [Warden AI](#). See the results of Beamery's independent bias audit [here](#).
- **Ongoing auditing:** Warden AI's auditing platform connects to Beamery's systems to provide ongoing monitoring for bias and fairness. The monthly audit frequency reflects the rapidly evolving nature of AI technology and Beamery's frequent development cycle. See live audit results [here](#).
- **Transparent:** The results of Warden's recurring audits are published publicly on Warden's independent website (see link above). This gives stakeholders clear oversight of our system and underscores our commitment to transparency.
- **Robust methodology:** Warden uses a comprehensive bias auditing approach with multiple techniques. It evaluates group fairness by examining outcome disparities between demographic groups, such as sex and race/ethnicity. Complementing this, Warden also assesses individual fairness to ensure similar individuals are treated consistently, analyzing outputs for profiles that are the same except for protected attributes. If no biases are detected using either technique, this combination provides reasonable assurance that the system is fair and trustworthy to use.
- **Compliant with AI regulations and standards:** Warden's approach meets and exceeds the requirements for bias auditing/testing stipulated by current and emerging AI regulations and standards, including NYC Local Law 144, the EU AI Act, ISO 42001, and the Colorado AI Consumer Protection Law. We're proud that the measures we have in place deliver a higher level of assurance than required by the legal minimum and will maintain compliance as new regulations come into effect.
- **Employer audits:** Beamery cannot guarantee that a user's specific use of the tool will never result in unfair bias (especially if the tools are not used as intended). For additional assurance, Beamery supports bias audits for individual employers using our platform. Employers can engage Warden to audit their use of Beamery's AI, providing an additional layer of assurance and compliance.

6. Social and Environmental Well-Being:

Creators of AI systems should be designing AI systems to contribute to sustainable and inclusive growth, social progress, and environmental well-being. Providers should consider the potential impact of AI systems on society and the environment.

We intend for our AI to have a meaningful impact on people. The Beamery system allows our customers to leverage AI for positive social benefits and to further equal access to meaningful work, skills, and careers for individuals from all backgrounds.

HOW DO OUR CUSTOMERS ACCOMPLISH THIS USING BEAMERY?

Transforming to [Skills-Based Hiring](#). A key benefit of skills-based hiring in your TA process will be an increase in diversity, equity, and inclusion. When skills are considered above everything else, the candidate pool becomes much larger, ensuring that everyone with the right skills is given a greater opportunity for recognition in the application process. When sourcing and hiring are carried out based on a standard and clearly defined skills language, and with the aid of automation or explainable AI helping to identify suitable candidates, bias towards traditional metrics of success (university, gender, subjective criteria) is reduced. This leads to fairer, more equitable shortlisting, and recruitment decisions can be made. Discover why a skills-first approach supports organizational diversity [here](#).

No system is perfect. These technologies involve a level of complexity necessitating a commitment to continual review and improvement as the field rapidly advances. We review new risk mitigation technologies as they become available and will consider employing them as appropriate.

Feedback: At Beamery, we ensure that feedback and perspectives are captured during the development process via:

- [Feedback Loops](#): Beamery has mechanisms in place to gather feedback about model behavior, and uses this feedback to improve.
- [External Reviews](#): External input and reviews are sought to ensure that no single internal perspective dominates.

Environmental Impact: At Beamery, we are mindful of the computing power required for AI, especially LLMs, which consume significant amounts of energy and can have a considerable effect on the environment. While we outsource our data center operations to responsible providers like Google, AWS, and Microsoft, Beamery actively seeks to reduce emissions within our own operations wherever possible. For more details, please refer to [Beamery's Impact & Sustainability Manifesto](#).

Part C

Responsible AI Procurement Considerations

In addition to the information presented above, we are making the following information available for our customers and prospective customers, so they have the information needed to follow the AI Procurement Best Practices created by the Center for Inclusive Change set out [here](#). Note: some of the principles and information below are duplicative of the above as the responsible AI procurement practices mirror those of the EU AI Act.

Beamery understands that making an algorithmic system explainable is key to building trust in such a system. Our aim is to maximize our end users' understanding of the following four AI explainability factors and how Beamery addresses them.

1. Transparency:

Logic & Rationale: How will (or did) the system arrive at the output?

We strive to ensure our customers understand each AI-driven recommendation, prediction, or output so they can make informed decisions, identify unintended outcomes, and mitigate any potential harm. There are numerous features within Beamery's services that are powered by AI, all of which can be summarized into providing three types of outputs:

- 1. Matches between candidates and vacancies** – Using Natural Language Processing (NLP), Deep Learning, and Graph Learning, Beamery's AI creates candidate and vacancy vector embeddings (both the vacancy as a profile of what an ideal candidate for that vacancy is). With this model, the output is the same representation for candidate and vacancy allowing the model to do the following matches: vacancy <-> vacancy; candidate <-> candidate and vacancy <-> candidate matching. This is then used by the end user to find most relevant candidates or vacancies for a given vacancy/candidates and provide a match score based on the weighting of five key factors: job title, skills, organization size, industry, seniority.
- 2. Inferred and suggested skills** – Using Deep Learning and Graph Learning, Beamery's AI can understand the relative association between entities of the same or different entity type, and recommend such connected entities to users. For example, it can be used to find an associated skill (one entity) to a role title (a different entity) based on patterns learned from our training data.
- 3. Generative AI output** – Using Microsoft Azure OpenAI LLMs, Beamery's TalentGPT can understand and generate text in a way that feels natural, drawing on its extensive training data to provide useful responses to questions, prompts or search requests. When you ask a question or give a prompt, GenAI uses the context of your input to generate a response, analyzing the sequence of words and identifying key concepts to predict an appropriate response.

Beamery's AI features are designed to be self-explanatory for new users; however, our customer-facing teams will be available to provide support about using AI features in addition to support articles available [here](#). The Beamery AI team can also provide additional documentation about model logic and evaluation that can be used in knowledge transfer sessions between Beamery and the Customer.



Data Sources & Uses: What data is/was used (and how) to create the output?

When training our models, we use de-identified third-party datasets consisting of over a billion data points. For more information, see above: Part B, Section 5 (Data Sources). We test our models using diverse and representative data sets that are most appropriate for the model's intended use. For example, we do not utilize data on trade union membership, as it does not serve the model's intended purpose.

2. Reliability:

Safety & Performance: How has the system been designed to ensure safety for its users?
(accuracy, robustness, reliability, and security)

All of Beamery's AI models have gone through several stages of learning and testing to improve their accuracy, using graph learning and representation learning with large samples of data (millions of profiles). All of our models also have specialized evaluation tasks that are designed to test different experiments against the 'real-life' use cases of the algorithm within the Beamery application.

Before adding a model into production, Beamery validates it with these benchmark metrics, and results/recommendations are spot-checked by employees at Beamery as part of our qualitative evaluation and quality assurance testing, to ensure that any model going into production has had multiple human reviews to capture any edge cases and limitations that might not have been considered by the quantitative evaluation.

However, Beamery has taken the responsible decision not to make any guarantees/warranties on accuracy as, in the context of Beamery's services, accuracy is relative, highly user-subjective, and difficult to measure. For example, one person's 5-star match may not be another person's 5-star match, even for the same vacancy. Further, it is well known that while machine learning/deep learning models can yield better "accuracy", the explainability of the outputs of such models is more difficult, which presents a challenge to model developers obligations of transparency.

When a model has more independent features or very few dependent features, it is more likely to be easily explainable (ie this input = this output). However, the model types used by Beamery, such as neural networks and XGBoost, has LOTS of dependent features by nature, which make them HIGHLY complex (less explainable in absolute terms. Lots of dependencies make the model output harder to explain in black-and-white terms (ie this input = this certain output). Lastly, it is important to take away that 100% accuracy is not possible for certain AI at this time.

For our robustness, reliability, and security operating processes, please see <https://trust.beamery.com/>. Our AI features undergo the same rigorous security auditing that all Beamery systems undergo.

3. Fairness:

Fairness: What steps have been taken to ensure fairness? (Ethical choices, data decisions, bias testing, output tests, user training, system monitoring, etc.)

- Beamery has robust bias auditing processes in place, including an independent third-party audit. For more information, see Part B, Section 5 (Diversity, Non-discrimination, and Fairness) above.
- Demographic data, such as sex, age, race, ethnicity, and sexual orientation, could have predictive value for recruitment cases. However, Beamery does not use such data in the training of any AI algorithms to prevent introducing unwanted bias into the process. Instead, we use a “pre-process” approach that abides by long-standing anti-discrimination laws by removing protected characteristics from training data as well as information or characteristics that could serve as proxies for protected characteristics, such as postal codes or neighborhoods known for certain ethnic populations.
- To further ensure fairness, we utilize the context of normalization. This is the process of ‘normalizing’ a free-form job title to a concept. For example, there are many ways to write “Java Developer”, ranging from “J2EE Full Stack Engineer” to “Java Developer/Engineer”, and everything in between. This example is a low-risk problem for AI as the result of the job title normalization is not influenced by a person’s protected characteristics only by the free-form job title. However, bias may arise through the properties of the language used. For example, a simple term-based search on “barman” in other systems may favor male candidates for that job, since the job title is male by definition. But, in Beamery, when recruiting for a barman job, the query will be automatically normalized (via our entity reconciliation algorithm) to remove gender bias inherent to that job title, to become: “bartender”. Everything is normalized to the core entities in our knowledge graph, and our algorithms act from there. A similar bias reduction can be achieved by normalizing the job titles (as discussed before), normalizing skills, and using it in queries: this ensures that no matter how the candidate expresses a skill or previous experience, the concept will still be matched.
- Beamery tests our algorithms for fairness at each key stage of the machine-learning process to detect and eliminate adverse impacts. Beamery utilizes internal fairness, accountability, transparency, and ethics (FATE) tools to check for a balanced output of results and determine if there is any bias against specific minority groups. These include various tests performed on training data, model training, and live model outcomes.

Impact: What steps have been taken to understand how the system will impact individuals and what was done to minimize negative, unfair, and/or adverse impacts? The list below represents a few of the desired critical impacts of Beamery’s AI to our customers, candidates, and greater society:

1. Increased efficiency when finding and evaluating candidates and reducing time to hire by suggesting or resurfacing contacts that are a good match for a particular vacancy;
2. Reduced bias in hiring and talent placement processes by enlarging the talent pools companies can choose talent from and focusing on elements that are relevant for the successful execution of the job (e.g. skills) rather than traditional bias-prone factors (e.g. the university an individual attended);
3. Improved transparency, objectivity, and understanding of how opportunities fit talent and further objectivity of how talent is matched for opportunities, reducing opacity in human decision-making on hires;
4. Through configured Job Architecture, we are enabling large organizations to avoid waste and regrettable attrition and/or redundancies of talent through a better understanding of their full talent profile and keeping employees engaged by deploying them where necessary in the business;
5. Rapid development and upskilling/reskilling of talent to meet organizations’ and society’s needs of the future.

This Explainability Statement serves as just one example of all the steps Beamery takes to ensure that our AI not only avoids negative impact on individuals, but also actively promotes positive outcomes for both individuals and society.



4. Accountability:

Responsibility: Who is responsible for the AI? (Who developed and deployed the system? Who is monitoring and addressing concerns, errors, and issues?)

- Beamery develops AI systems (a “System Provider” under the EU AI Act), while our customers deploy them (a “Deployer” under the EU AI Act). We proactively monitor our systems internally and address any concerns, errors, or other issues as they arise. Additionally, through coordinated efforts and open dialogue with our customers, we encourage and solicit feedback from users of the system.
- Beamery educates and empowers our customers to make informed decisions about using AI responsibly. We achieve this by providing tools and resources that help our customers better understand the system and its potential impacts, enabling them to understand their responsibility in adopting AI in a safe and reliable way.
- All Beamery AI features are deployed in a way that can be easily switched on or off at will, so customers can opt out of AI usage at any time. In addition to this, Beamery’s Talent Match Scoring can be disabled for applicants (and used only for Suggested Contacts). Despite this option to forgo our AI, our clients – including some of the world’s largest companies – have chosen to leverage Beamery’s AI functionality, highlighting the significant value it adds to their Beamery experience.

DISCLAIMER: This document is for informational purposes only. This document is not intended to, and does not, provide legal advice or a legal opinion. This document may be updated from time to time based on updates to our internal processes or products or updates in applicable law and guidance.



How does Beamery use LLMs to prove humanlike interactions across the lifecycle?



Beamery TalentGPT

Our interactive experience allows users to ask questions in a natural way



LLM Microsoft Azure

Third party that interprets the question in language form and translates the “ask” into a query Beamery’s AI can understand



Beamery AI

Beamery’s proprietary, explainable and audited AI model that queries our data to provide an answer to the question



LLM Microsoft Azure

Third party that takes in this data from Beamery AI and translates it back into a language form to articulate a response



Beamery TalentGPT

Our interactive experience that allows users to ask questions in a natural way

figure 2: Diagram depicting how Beamery’s Gen AI models use Microsoft Azure’s OpenAI LMM in connection with the above models.